arXiv:1612.04668v2 [physics.soc-ph] 17 Dec 2016

Multi-resolution community detection in massive networks

Jihui Han[1,*], Wei Li[1,+], and Weibing Deng[1,@]

[1]Complexity Science Center and Institute of Particle Physics, Central China Normal University, Wuhan, 430079, China

[*]jh@mails.ccnu.edu.cn

[+]liw@mail.ccnu.edu.cn

[@]liw@mail.ccnu.edu.cn

ABSTRACT

Aiming at improving the efficiency and accuracy of community detection in complex networks, we proposed a new algorithm, which is based on the idea that communities could be detected from subnetworks by comparing the internal and external cohesion of each subnetwork. In our method, similar nodes are firstly gathered into meta-communities, which are then decided to be retained or merged through a multilevel label propagation process, until all of them meet our community criterion. Our algorithm requires neither any priori information of communities nor optimization of any objective function. Experimental results on both synthetic and real-world networks show that, our algorithm performs quite well and runs extremely fast, compared with several other popular algorithms. By tuning a resolution parameter, we can also observe communities at different scales, so this could reveal the hierarchical structure of the network. To further explore the effectiveness of our method, we applied it to the E-Coli transcriptional regulatory network, and found that all the identified modules have strong structural and functional coherence.

Introduction

A wide range of natural and social systems can be described as complex networks[1–6]. Examples include the cell, a network of chemicals linked by chemical reactions, or the Internet, a network of routers and computers connected by physical links.

Most real-world networks are observed to contain communities[7]. Intuitively, a community is a group of nodes which are relatively densely connected to each other within the group but sparsely connected to the nodes in other groups of the network[8]. Detecting communities can not only uncover the relations between internal structures and functional behaviours of networks, but also have many practical applications in the domains such as biology, sociology, economics and computer science[9–11]. Therefore, it is not surprising that community detection has been so extensively investigated during the past few years[7].

In the last decade, lots of methods have been developed to detect the community structure, such as modularity optimization[12–16], dynamic label propagation[17–21], statistical inference[22,23], Potts model[24,25], spectral clustering[26,27], information-

theoretic[28,29] and topology based[30–33] methods. Among them, one popular group of approaches are based on the optimization of modularity, which could be more or less subject to the resolution problem[34]. Especially fast ones are label propagation algorithms, in which each node adopts the majority label among its neighbours, and the labels propagate iteratively until there is no change in the network. Due to the frequent tie-breaks and random processing order of nodes, label propagation algorithms usually deliver multiple partitions starting from the same initial condition with different random seeds.

Among all the community detection methods developed so far, it is yet unlikely to obtain a widely accepted definition of community. Hence, we first introduce a quantitative criterion to determine what kind of subnetworks are communities, and then propose a fast algorithm to detect them. The proposed method mainly consists of two steps: initialization and multilevel label propagation. The initialization step is implemented to form some meta-communities by collecting similar nodes. While the multilevel label propagation step, which is similar to the Louvain method[16], consists of two sub-steps: network collapse and label propagation. Firstly, it aggregates nodes that belong to the same meta-community and builds a new network whose nodes represent the meta-communities detected in the previous step. Secondly, it retains or merges meta-communities by comparing their internal and external connections (or weights) through a weighted version of label propagation. These two sub-steps repeat iteratively until all the meta-communities meet our community criterion. As discussed above, our method requires neither optimization of objective function nor any prior information of communities. We tested our algorithm on both synthetic and real-world networks, and compared it with several other popular algorithms. Results show that our algorithm could detect meaningful communities in large networks efficiently and accurately, and it could also uncover the hierarchical structure of the network by tuning a resolution parameter.

Results

We evaluated the performance of our method on both synthetic and real-world networks. For synthetic networks, we tested the classical benchmark proposed by Girvan and Newman (GN)[35], and the well-known benchmark with planted community structure and heterogeneous distributions of node degree and community size proposed by Lancichinetti, Fortunato and Radicchi (LFR)[36]. As real-world networks have some different topological properties that distinguish them from the synthetic ones, we also tested our method on different kinds of real-world networks, such as social networks and biological networks. To assess its performance, we compared our algorithm with other six popular algorithms listed in Table 1, in terms of normalized mutual information (NMI) and modularity.

Synthetic Networks

In this section, we tested our method against synthetic benchmarks. In our algorithm, we set $\lambda = 1$ by default (λ can be regarded as a resolution parameter that controls the scale on which we would like to observe the communities in a network, see **Methods** for detailed discussions). In all tests on synthetic networks, each point is always an average over 100 different network realizations. We adopted NMI as a measure of consistency between the planted partition and the detected one.

Tests on the GN benchmark

We first tested our method on the GN benchmark, which is the most famous benchmark for community detection. The GN network consists of 128 nodes which are divided into four equal groups. The edges are placed independently and randomly between node pairs, with probability p_{in} for edges to fall between nodes in the same community, and p_{out} for edges to fall between nodes in different communities. The values of p_{in} and p_{out} are chosen to make the expected degree of each node equal 16, and thus not independent. Here, we choose the mixing parameter μ as an independent parameter, which indicates the ratio of the external degree of a node with respect to its community to the total degree of the node.

Figure 1 (a) shows the average NMI scores of different algorithms. As can be seen, LP fails to detect the communities even for small μ ($\mu \sim 0.3$). Though better than LP, LE does not have a remarkable performance either, as it also starts to fail for low values of μ. Our algorithm and Infomap have comparable performance. Both of them are better than LE, but outperformed by modularity-based methods: GN, FADM and the Louvain method. GN performs nearly as well as the Louvain method does. FADM performs the best, which yields a high NMI up until $\mu \sim 0.5$, where the rest algorithms cannot detect communities accurately.

Figure 1 (b) shows the execution time of different algorithms as a function of μ. As we can see, our algorithm runs extremely fast, even faster than LP. This could be that the consideration of weights (or similarities) reduces the label oscillations compared with LP, which enables our algorithm to converge faster. The Louvain method is sightly slower than LP, but significantly faster than the rest of four methods. Infomap and FADM have comparable execution time, and both of them are sightly slower than LE. GN is the slowest among all algorithms.

Tests on the LFR benchmark

We then tested our method on the LFR benchmark, and compared the results with those of the other methods. In the LFR network, both the degree of each node and the size of each community are drawn from power-law distributions. Each node shares on average a fraction $1-\mu$ of its edges with the other nodes of its community and a fraction μ with the nodes of the other communities; $0 \leq \mu \leq 1$ is the mixing parameter which indicates the significance of community structure.

In Figure 2, we show the average NMI and execution time of different algorithms as a function of the mixing parameter on the LFR networks. The following parameters apply to all LFR networks used here: the average degree and the maximum degree are 20 and 50 respectively, the exponents of the degree distribution and the community size distribution are -2 and -1 respectively. We show four sample results which correspond to two different network sizes (1000 and 5000), and two different ranges regarding community sizes, indicated by the letters 'S' (small communities with 10 to 50 nodes) and 'B' (big communities with 20 to 100 nodes), respectively. For the GN algorithm, we only show the results on smaller networks due to the high computational complexity of the method.

Generally, modularity-based methods, such as GN, LE, FADM and the Louvain method, have rather poor performance, especially in the case of larger networks with smaller communities, due to the well-known resolution limit of modularity[34]. LP does not have impressive performance either and its performance is sightly affected by the size of communities, i.e., it performs slightly better in the case of smaller communities than in the case of larger communities. Infomap performs sightly better than

our algorithm on larger networks, and has comparable performance with our algorithm in the case of smaller networks and communities. However, it is outperformed by our method in the case of smaller networks with bigger communities. The results confirm that our algorithm has a reasonable performance on the LFR networks. Moreover, our algorithm runs extremely fast and is only sightly slower than LP (see Figure 2 (b)).

Tests on random networks

We also tested our method on random networks. In random networks, the connecting probabilities of the nodes are independent of each other, which leads to homogeneous density of links in the network. Thus, there should be no community structures. A good algorithm should not find non-trivial partitions and ideally deliver only one community which contains all nodes of the network.

Here, we considered two types of random networks: Erdős-Rényi[37] (ER) and scale-free[38] (SF). In ER random networks, nodes have the same probability to get connected to each other and the degree distribution is then binomial. The SF random networks are built via the configuration model[39], by starting from a certain degree sequence for the nodes obeying the predefined power law distribution with exponent -2. The sizes of all networks are fixed to 1000.

Figure 3 shows the number of communities detected by different methods as a function of the average degree. The GN algorithm is excluded because it is too slow to be used for analysis. In both ER and SF random networks, LP, Infomap and our method always find a single community containing all nodes of the network except when the average degree is small. However, modularity-based methods, such as LE, FADM and the Louvain method, are not so good, as they always find a few communities even for a large average degree. These results show that our algorithm tends to find a few small communities in a sparse random network (due to stochastic fluctuations, specific realizations of random networks may display pseudo-communities), while it always detects a single community in a random network with dense connections.

Analysis of resolution limit

Finally, we analysed a kind of network made of some identical complete networks (cliques), connected by a single edge with each other, and each clique contains three nodes at least. Ideally, a good algorithm should detect all the predefined cliques. Figure 4 shows that our method detects all the cliques in all these examples. Infomap finds all 4-node cliques, but fails to detect all the 3-node cliques when the network gets large (i.e., with more than 26 cliques). LP finds sightly less communities than the predefined ones. The remaining four modularity-based methods detect dramatically less communities than the predefined ones when the network is large enough. This is mainly due to the resolution limit of modularity[34], i.e., the optimization of modularity may fail to identify small communities below a cutoff size depending on the network size.

Real-World Networks

The real-world networks are far more complex than the synthetic ones. Thus, it is still a great challenge to uncover the community structure of real-world networks. In this section, we applied our method to 9 different real-world networks listed in Table 2. The sizes of these networks range from tens to millions. We presented the detailed results as follows.

Zachary's karate club is a social network of friendships between 34 members of a karate club at a US university in the

1970s. It was divided into two smaller clubs after a dispute between club president John (node 34) and instructor Mr. Hi (node 1). When $\lambda = 0.6$, two communities are detected by our algorithm, which are identical to the two real ones, as shown in Figure 5 (a). When $\lambda = 1$, one of the two communities is divided into two smaller ones, as shown in Figure 5 (b).

The dolphin social network describes the frequent associations between 62 dolphins living off Doubtful Sound, New Zealand. The links represent that dolphins are observed to stay together more often than expected by chance during the years from 1994 to 2001. Figure 5 (c) shows the two communities identified by our algorithm with $\lambda = 0.6$, which are identical to the two real ones except for the node 'SN89'. When λ increases to 1, one of the two communities is divided into three smaller ones, as shown in Figure 5 (d).

Books about US politics network is compiled by Valdis Krebs. Nodes represent books about US politics sold by the online bookseller Amazon.com. Edges represent frequent co-purchasing of books by the same buyers. Books have been divided into liberal, neutral, or conservative with respect to their attitudes. As shown in Fig. 6 (c), our algorithm identifies three communities which resemble the three natural ones.

The American college football network describes football games among Division IA colleges during regular season Fall 2000. As shown in Fig. 6 (a), 115 nodes in the network represent teams (identified by their college names), which are grouped into eleven different conferences, except for five independent teams (Utah State, Navy, Notre Dame, Connecticut and Central Florida). The regular season games between each pair of teams are shown as 613 edges of the network. When $\lambda = 0.6$, our algorithm identifies elven communities within this network. Among them, eight conferences (i.e., Atlantic Coast, Big East, Big Ten, Big Twelve, Mid-American, Mountain West, Pacific Ten and Southeastern) are correctly identified. The three remaining communities closely resemble the Conference USA, Sun Belt and Western Athletic conferences. Five independent teams that do not belong to any conference tend to be grouped with the conferences which are most closely associated.

The E. Coli transcriptional regulatory network is a directed biological network with 578 edges and 423 nodes which is compiled by Shen-Orr et al. in 2002. Nodes are operons and edges start from an operon that encodes a transcription factor to another operon with which it directly regulates. Here we used an undirected version of the network described in the updated RegulonDB[40]. We identified 5 isolated nodes, 26 modules with two operons, 9 modules with three operons and 22 modules with more than three operons (see Supplementary Fig. S9 online). We analysed the 22 modules that have more than 3 elements with the DAVID functional annotation tool[41,42]. The greater probability that the genes appear to participate in a common biological process is described with smaller p-values. As shown in Table 3, all these modules are functionally coherent. For example, the first module contains 53 operons, which performs as the cellular respiration (p-value is 5.2E-89). The second module has 6 operons and is involved in "aromatic amino acid family biosynthesis process" (p-value is 3.2E-13). The results of other modules are listed in Table 3. The entire operon list of the 22 modules can be found in Supplementary Table S2 online.

The *Polblogs* network is a directed network of hyperlinks between weblogs on US politics, recorded in 2005 by Adamic and Glance. Here we use an undirected version of the network. As shown in Fig. 6 (b), the resulting partition obtained by our algorithm mainly contains 2 communities, which indicates the liberal and conservative political leaning.

The *Facebook* network describes friendships between users, which was collected from survey participants by using Face-

book application. There are 4039 users and 88218 friendships in this network. We detected 7 communities by using our algorithm with $\lambda = 0.05$ and 19 communities with $\lambda = 0.2$, as shown in Fig. 5 (e) and (f), respectively. The community structure of this network is quite clear based on visual observation.

The *Amazon* network is collected by crawling the Amazon website. It is based on "Customers Who Bought This Item Also Bought" feature of the Amazon website. If a product is frequently co-purchased with another product, the network contains an undirected edge between these two products. The Youtube social network contains the friendships of users. We analysed these two networks by using our algorithm with $\lambda = 1$ and presented the distribution of community sizes in Figure 7. It is found that the community sizes of both Amazon and Youtube networks can be well fitted by the power-law distributions, with exponents -2.83 ($x_{min} = 10$, p-value=0.49) and -2.61 ($x_{min} = 5$, p-value=0.45) respectively[43]. This is consistent with the previous observations on other social networks[14,17,44].

We also compared the results with those of the other methods, and showed their performances in Figure 8. The results of GN, LE, and FADM on some large networks are absent due to high computational cost. As can be seen, the composite scores (i.e., the sum of NMI and modularity) of our algorithm are competitive, especially on the *Karate* and *Dolphins* networks. Moreover, our algorithm runs extremely fast, even faster than LP on most networks. Therefore, our algorithm can detect communities on very large networks in short time with a reasonable accuracy.

Analysis of the resolution parameter

In this subsection, we investigated the resolution parameter λ in our method by studying four real-world networks, which are *Karate*, *Dolphins*, *Polbooks* and *Football* networks. As shown in Figure 9 (c), the number of detected communities will increase when λ gets large. This is because the parameter λ controls the weight of cohesion within a subnetwork. With the increase of λ, the importance of the cohesion will increase, and communities tend to be small and tight. While, when λ is small, especially equal to 0, the attractions between subnetworks are always larger than their cohesions, thus the whole network will become one community eventually. In particular, the larger the resolution parameter, the smaller the detected communities.

In Figures 9 (a) and (b), we show the NMI and modularity of community structures at different scales. In general, the trend of community qualities for each network is like a stair, which increases from a low quality to a high value through some plateaus. This indicates the parameter λ has some stable ranges, which could correspond to community structures of the network at different scales. In order to reveal the hierarchical structure of the network, we can employ a large resolution parameter firstly (e.g., $\lambda = 1$), to obtain the community structure at low scales. Then, by decreasing the resolution parameter gradually, we can explore the community structures at high scales. Finally, the hierarchical structure of the network could be unfolded.

Discussion

We propose a new method to detect community structures in complex networks. In our approach, similar nodes are first grouped together into meta-communities which will be retained or merged through a multilevel label propagation process

(see **Methods** for details). We introduce a quantitative community criterion to determine what kind of subnetworks are communities. With the aid of this criterion, our algorithm can stop at an appropriate level when all the subnetworks meet our community criterion, and thus determine the number of communities automatically. Moreover, by tuning a resolution parameter, we can reveal the hierarchical organization of the network.

Compared with several other popular algorithms, our method has robust performance (in terms of NMI) on both synthetic and real-world networks with which ground truth is known, and the modularity scores are also competitive in most of the real-world networks. Moreover, our algorithm has a linear time complexity and runs extremely fast, which enables it to handle very large networks.

In our experiments, we observe that the modularity obtained by our method shows an increasing trend during the label diffusion process (see Supplementary Fig. S3 online), although the maximization of modularity is not our intention. This provides further experimental evidence for the effectiveness of our method. In contrast to modularity-based methods, according to the testing results, our algorithm doesn't need to merge small communities to have a higher modularity, and thus preserves them even in large networks. Therefore, our method eliminates the resolution limit of modularity-based methods.

Moreover, our algorithm generally requires only a small number of iterations even on large networks (see Supplementary Fig. S2 online). The number of meta-communities decreases dramatically in each iteration. Thus, most computing time is consumed in the calculation of similarities and the first iteration. In practice, if we want to explore the hierarchical structure of a network, we only need to calculate the similarities once, and then perform our algorithm on the network with different values of the resolution parameter. This makes our method faster in analysing the hierarchical organization of a network.

Methods

The key intuition behind our method is that, similar and tightly connected nodes are likely to belong to the same community. In the following subsections, we first define the similarity between two adjacent nodes, and introduce a criterion to determine what kind of subnetworks are communities. Then we present our algorithm in detail.

Structural similarity

There are different measures based on the common neighbourhood to quantify the similarity of any pair of adjacent nodes. Here, we use the structural similarity measure given by the cosine similarity function[45], which effectively denotes the local connectivity density of any two adjacent nodes in a weighted network. Given an undirected weighted network $G = (V, E, w)$, the structural similarity $s(u,v)$ between two adjacent nodes u and v is defined as:

$$s(u,v) = \frac{\sum_{x \in \Gamma(u) \cap \Gamma(v)} w(u,x) \cdot w(v,x)}{\sqrt{\sum_{x \in \Gamma(u)} w^2(u,x)} \cdot \sqrt{\sum_{x \in \Gamma(v)} w^2(v,x)}}, \tag{1}$$

where $\Gamma(u) = \{v \in V | \{u,v\} \in E\} \cup \{u\}$ is the set containing u and its adjacent nodes. If the network is unweighted, equation (1) can be simplified to be,

$$s(u,v) = \frac{|\Gamma(u) \cap \Gamma(v)|}{\sqrt{|\Gamma(u)| \cdot |\Gamma(v)|}}. \tag{2}$$

$s(u,v)$ corresponds to the so-called edge-clustering coefficient introduced by Radicchi et al[46]. The similarity value is in the range $(0,1]$. Two nodes will have a higher similarity by sharing more common neighbours. Once the neighbours of the two nodes are exactly the same, the similarity measure equals 1.

Community criterion

We define a subnetwork as a community by comparing its internal and external connections (or weights). An undirected unweighted network G can be represented by an adjacency matrix A with entries $A_{uv} = 1$ if u is directly connected to v and $A_{uv} = 0$ otherwise. The subnetwork V_i is a community if for any other subnetwork V_j

$$\lambda \sum_{u \in V_i, v \in V_i} A_{uv} \geq \sum_{u \in V_i, v \in V_j} A_{uv}. \tag{3}$$

Informally, a subnetwork is a community if its internal connections exceeds the number of edges shared by the subnetwork with the other communities. Our definition of community is in the same spirit of weak community proposed by Hu et al[47]. The difference is that we introduce a parameter λ to quantitatively compare the internal and external connections (or weights) of each subnetwork. We can deem that $\lambda \sum_{u \in V_i, v \in V_i} A_{uv}$ is the cohesion of subnetwork V_i, and $\sum_{u \in V_i, v \in V_j} A_{uv}$ is the attraction between subnetworks V_i and V_j. If λ is very small, subnetworks tend to be attracted to each other and are merged to form large communities. On the contrary, when λ is close to 1, subnetworks tend to form communities by their own. If λ equals 1, the definition of weak community proposed by Hu et al.[47] is recovered. Therefore, parameter λ can be regarded as a resolution parameter which controls the scale upon which we would like to observe the communities in a network. Large λ yields small, tight communities, and small λ instead reveals large communities. In most situations, the whole network forms a single community when λ approaches 0. In contrast, natural communities which are commonly detected by other algorithms are identified when $\lambda = 1$.

Algorithm

Our algorithm is based on the idea that communities are groups of nodes which are similar to each other, and communities can be detected from subnetworks by comparing the internal and external connections (or weights) of each subnetwork. The algorithm mainly consists of two steps: initialization and multilevel label propagation.

Initialization. It is observed that, the more similar two nodes are, the more likely they are assigned to the same community. So the similarity measure based on common neighbours provides a natural way to detect communities in networks. Therefore, in the initialization step, we first calculate the similarity between each pair of adjacent nodes. Then we re-assign labels to nodes in a way that each node takes the most similar label of its neighbours in a asynchronous manner. This label propagation process proceeds iteratively until the label of each node is one of the most similar labels in its neighbourhood. After the propagation process, we should obtain some meta-communities which consist of most similar nodes.

Multilevel label propagation. In general, the meta-communities detected in the initialization step do not all meet our community criterion, so we need the multilevel label propagation step to merge some of them. This step consists of two sub-steps: network collapse and label propagation, which are repeated iteratively.

Firstly, we built a new network whose nodes represent the meta-communities that were detected in the previous step. The weights of edges between the new nodes are given by the total weight of the edges between nodes inside the corresponding meta-communities. The total weight between nodes inside the same meta-community leads the weight of self-loop for this new node in the collapsed network. Once this sub-step is finished, we obtain a weighted collapsed network with weights representing the cohesion of meta-communities or attraction between meta-communities in the un-collapsed network.

Secondly, we perform a label propagation process on the collapsed network. To detect communities that meet the community criterion, we rescale the weights of self-loops by λ prior to the label propagation process. Essentially, during the label propagation process, the meta-communities (i.e., the macro-nodes in the collapsed network) compare their internal cohesion with external attraction, and then decide to be retained or merged. After the label propagation, we should obtain some meta-communities of macro-nodes. Then it is possible to re-apply the first sub-step to collapse the network.

The number of meta-communities decreases dramatically in each iteration, and as a consequence most of the computing time is consumed in the first iteration. These two sub-steps are iterated until all the meta-communities meet our community criterion, i.e., the cohesion of each meta-community is greater than the attraction between it and any other meta-community. That is, meta-communities are no longer merged and thus the number of meta-communities no longer changes. Generally, the algorithm requires a small number of iterations (see Supplementary Fig. S2 online). The detailed description of the algorithm and an illustration of the application of the algorithm to the *Karate* network are shown in the supplementary material.

Evaluation measures

To evaluate the effectiveness of the proposed algorithm we use the following two quality measures: modularity and normalized mutual information. As a widely used quality measure, modularity is defined by[12]:

$$Q = \sum \left(e_{ii} - a_i^2 \right), \tag{4}$$

where e_{ij} represents the fraction of total connections between two different communities, e_{ii} represents the real fraction of links exist within a community, $a_i = \sum_j e_{ij}$ corresponds to the fraction of links connected to community i, and the expected number of intra-community links is just a_i^2. Modularity is based on the intuitive idea that random networks do not exhibit community structure.

For a network with known community structure, consistency of the detected and the true partition is quantified by using the normalized mutual information[48]:

$$I(A,B) = \frac{-2\sum_{i=1}^{c_A}\sum_{j=1}^{c_B} N_{ij}\log(N_{ij}N/N_{i.}N_{.j})}{\sum_{i=1}^{c_A} N_{i.}\log(N_{i.}/N) + \sum_{j=1}^{c_B} N_{.j}\log(N_{.j}/N)}, \tag{5}$$

where A and B represent the real and the detected partitions respectively. The number of real communities is denoted by c_A and the number of detected communities is denoted by c_B. N_{ij} is the number of nodes in the real community i that appear in the detected community j. The sum over row i of matrix N_{ij} is denoted by $N_{i.}$ and the sum over column j is denoted by $N_{.j}$. If the detected communities are identical to the real ones, $I(A,B)$ takes the maximum value 1. Whereas $I(A,B) = 0$ if the partition detected by the algorithm is totally independent of the real partition.

Time complexity

Given a network with n nodes and m edges. Let k be the maximum degree of nodes in this network. The time complexity of each step of our algorithm is roughly estimated as follows.

1. Similarity calculation takes time of $O(km)$ at most. For each pair of adjacent nodes, we need to iterate through all neighbours of the two nodes to calculate the similarity, and the upper bound of time complexity is $O(k)$. Thus calculating similarities for all pairs of adjacent nodes takes time of $O(km)$ at most.

2. Aggregating network needs to iterate through each edge of the network, which takes time of $O(m)$ at most.

3. Label propagation needs to iterate all nodes of the network, and each node iterates through at most k neighbours, thus the upper bound of time complexity is $O(kn)$.

Steps 2 and 3 repeat iteratively until all the communities meet the community criterion. Generally, our algorithm converges after a small number of iterations (see Supplementary Fig. S2 online), so the time complexity is roughly $O(kn+km)$. If the network is sparse (i.e., $m \sim n$) and $k \ll n$, our algorithm has a linear time complexity, and thus can be efficiently applied to large-scale networks.

References

1. Rubinov, M. & Sporns, O. Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* **52**, 1059 – 1069 (2010).

2. Broder, A. *et al.* Graph structure in the web. *Computer Networks* **33**, 309 – 320 (2000).

3. Palla, G., Barabási, A.-L. & Vicsek, T. Quantifying social group evolution. *Nature* **446**, 664–667 (2007).

4. Onnela, J.-P. *et al.* Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences* **104**, 7332–7336 (2007).

5. Barabási, A.-L. & Oltvai, Z. N. Network biology: understanding the cell's functional organization. *Nat Rev Genet* **5**, 101–113 (2004).

6. Ravasz, E., Somera, A. L., Mongru, D. A., Oltvai, Z. N. & Barabási, A.-L. Hierarchical organization of modularity in metabolic networks. *Science* **297**, 1551–1555 (2002).

7. Fortunato, S. Community detection in graphs. *Physics Reports* **486**, 75 – 174 (2010).

8. Porter, M. A., Onnela, J.-P. & Mucha, P. J. Communities in networks. *Notices of the AMS* **56**, 1082–1097 (2009).

9. Fan, M., Wong, K.-C., Ryu, T., Ravasi, T. & Gao, X. Secom: A novel hash seed and community detection based-approach for genome-scale protein domain identification. *PLoS ONE* **7**, 1–11 (2012).

10. Ratti, C. *et al.* Redrawing the map of great britain from a network of human interactions. *PLoS ONE* **5**, 1–6 (2010).

11. Tang, C., Li, X., Cao, L. & Zhan, J. The law of evolutionary dynamics in community-structured population. *Journal of Theoretical Biology* **306**, 1 – 6 (2012).

12. Newman, M. E. J. & Girvan, M. Finding and evaluating community structure in networks. *Phys. Rev. E* **69**, 026113 (2004).

13. Newman, M. E. J. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* **103**, 8577–8582 (2006).

14. Clauset, A., Newman, M. E. J. & Moore, C. Finding community structure in very large networks. *Phys. Rev. E* **70**, 066111 (2004).

15. Duch, J. & Arenas, A. Community detection in complex networks using extremal optimization. *Phys. Rev. E* **72**, 027104 (2005).

16. Blondel, V. D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **2008**, P10008 (2008).

17. Raghavan, U. N., Albert, R. & Kumara, S. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E* **76**, 036106 (2007).

18. Gregory, S. Finding overlapping communities in networks by label propagation. *New Journal of Physics* **12**, 103018 (2010).

19. Xie, J. & Szymanski, B. K. Labelrank: A stabilized label propagation algorithm for community detection in networks. In *Network Science Workshop (NSW), 2013 IEEE 2nd*, 138–143 (2013).

20. Lin, Z., Zheng, X., Xin, N. & Chen, D. Ck-lpa: Efficient community detection algorithm based on label propagation with community kernel. *Physica A: Statistical Mechanics and its Applications* **416**, 386 – 399 (2014).

21. Liu, W., Jiang, X., Pellegrini, M. & Wang, X. Discovering communities in complex networks by edge label propagation. *Scientific Reports* **6**, 22470 EP – (2016).

22. Karrer, B. & Newman, M. E. J. Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83**, 016107 (2011).

23. Aldecoa, R. & Marín, I. Surprise maximization reveals the community structure of complex networks. *Scientific Reports* **3**, 1060 EP – (2013).

24. Ronhovde, P. & Nussinov, Z. Multiresolution community detection for megascale networks by information-based replica correlations. *Phys. Rev. E* **80**, 016109 (2009). URL http://link.aps.org/doi/10.1103/PhysRevE.80.016109.

25. Ronhovde, P. & Nussinov, Z. Local resolution-limit-free potts model for community detection. *Phys. Rev. E* **81**, 046114 (2010). URL http://link.aps.org/doi/10.1103/PhysRevE.81.046114.

26. Newman, M. E. J. Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* **74**, 036104 (2006).

27. Singh, A. & Humphries, M. D. Finding communities in sparse networks. *Scientific Reports* **5**, 8828 EP – (2015).

28. Rosvall, M. & Bergstrom, C. T. An information-theoretic framework for resolving community structure in complex networks. *Proceedings of the National Academy of Sciences* **104**, 7327–7331 (2007).

29. Rosvall, M. & Bergstrom, C. T. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences* **105**, 1118–1123 (2008).

30. Liu, W., Pellegrini, M. & Wang, X. Detecting communities based on network topology. *Scientific Reports* **4**, 5739 EP – (2014).

31. Žalik, K. R. Maximal neighbor similarity reveals real communities in networks. *Scientific Reports* **5**, 18374 EP – (2015).

32. Chen, Y., Zhao, P., Li, P., Zhang, K. & Zhang, J. Finding communities by their centers. *Scientific Reports* **6**, 24017 EP – (2016).

33. Bagrow, J. P. & Bollt, E. M. Local method for detecting communities. *Phys. Rev. E* **72**, 046108 (2005).

34. Fortunato, S. & Barthélemy, M. Resolution limit in community detection. *Proceedings of the National Academy of Sciences* **104**, 36–41 (2007).

35. Girvan, M. & Newman, M. E. J. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* **99**, 7821–7826 (2002).

36. Lancichinetti, A., Fortunato, S. & Radicchi, F. Benchmark graphs for testing community detection algorithms. *Phys. Rev. E* **78**, 046110 (2008).

37. Erdős, P. & Rényi, A. On random graphs i. *Publ. Math. Debrecen* **6**, 290–297 (1959).

38. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).

39. Molloy, M. & Reed, B. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms* **6**, 161–180 (1995).

40. Gama-Castro, S. *et al.* Regulondb version 9.0: high-level integration of gene regulation, coexpression, motif clustering and beyond. *Nucleic Acids Res* **44(D1)**, D133–43 (2016).

41. Huang, D. W., Sherman, B. T. & Lempicki, R. A. Systematic and integrative analysis of large gene lists using david bioinformatics resources. *Nat. Protocols* **4**, 44–57 (2008).

42. Huang, D. W., Sherman, B. T. & Lempicki, R. A. Bioinformatics enrichment tools: paths toward the comprehensive functional analysis of large gene lists. *Nucleic Acids Research* **37**, 1–13 (2009).

43. Clauset, A., Shalizi, C. R. & Newman, M. E. J. Power-law distributions in empirical data. *SIAM Review* **51**, 661–703 (2009).

44. Newman, M. E. J. Fast algorithm for detecting community structure in networks. *Phys. Rev. E* **69**, 066133 (2004).

45. Huang, J., Sun, H., Han, J. & Feng, B. Density-based shrinkage for revealing hierarchical and overlapping community structure in networks. *Physica A: Statistical Mechanics and its Applications* **390**, 2160 – 2171 (2011).

46. Radicchi, F., Castellano, C., Cecconi, F., Loreto, V. & Parisi, D. Defining and identifying communities in networks. *Proc Natl Acad Sci U S A* **101**, 2658–2663 (2004).

47. Hu, Y. *et al.* Comparative definition of community and corresponding identifying algorithm. *Phys. Rev. E* **78**, 026121 (2008).

48. Danon, L., Díaz-Guilera, A., Duch, J. & Arenas, A. Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment* **2005**, P09008 (2005).

49. Csardi, G. & Nepusz, T. The igraph software package for complex network research. *InterJournal* **Complex Systems**, 1695 (2006). URL http://igraph.org.

50. III, S. T., Nyberg, A., Genio, C. I. D. & Bassler, K. E. Fast and accurate determination of modularity and its effect size. *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P02003 (2015).

51. Zachary, W. W. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* **33**, 452–473 (1977).

52. Lusseau, D. *et al.* The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations. *Behavioral Ecology and Sociobiology* **54**, 396–405 (2003).

53. Krebs, V. A network of co-purchased books about us politics sold by the online bookseller amazon.com (2008). URL http://www-personal.umich.edu/~mejn/netdata.

54. Newman, M. E. J. The structure and function of complex networks. *SIAM Review* **45**, 167–256 (2003).

55. Shen-Orr, S. S., Milo, R., Mangan, S. & Alon, U. Network motifs in the transcriptional regulation network of escherichia coli. *Nat Genet* **31**, 64–68 (2002).

56. Adamic, L. A. & Glance, N. The political blogosphere and the 2004 u.s. election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery*, LinkKDD '05, 36–43 (ACM, New York, NY, USA, 2005).

57. Leskovec, J. & Mcauley, J. J. Learning to discover social circles in ego networks. In Bartlett, P., Pereira, F., Burges, C., Bottou, L. & Weinberger, K. (eds.) *Advances in Neural Information Processing Systems 25*, 548–556 (2012).

58. Yang, J. & Leskovec, J. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems* **42**, 181–213 (2013).

59. Leskovec, J. & Krevl, A. SNAP Datasets: Stanford large network dataset collection. http://snap.stanford.edu/data (2014).

Acknowledgements

This work was in part supported by the Program of Introducing Talents of Discipline to Universities under grant no. B08033, and National Natural Science Foundation of China (Grant No. 11505071).

Author contributions statement

J.H. designed the algorithm, carried out the simulations and prepared all the tables and figures. J.H., W.L. and W.D. prepared the main manuscript text. All authors reviewed the manuscript.

Additional information

Competing financial interests: The authors declare no competing financial interests.

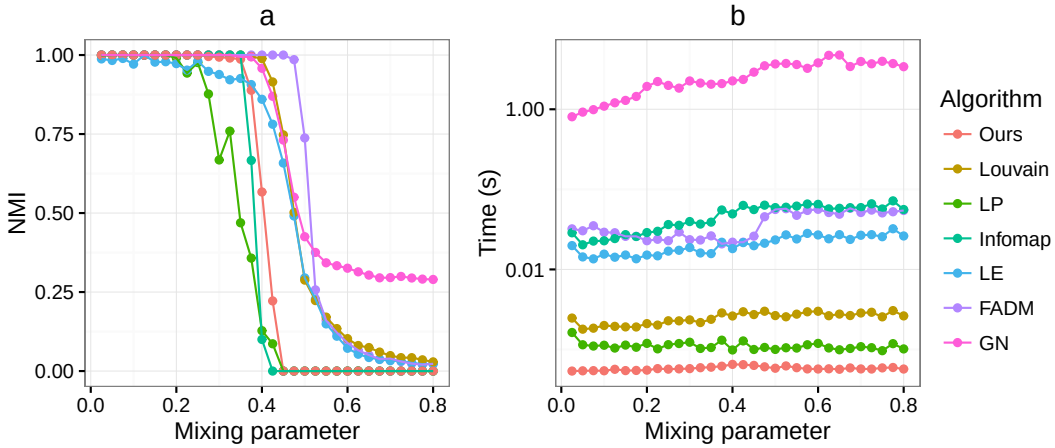

Figure 1. (Colour online) **Tests of the algorithms on the GN benchmark.** (a) shows the normalized mutual information as a function of the mixing parameter. (b) shows the execution time of different algorithms on the benchmark.

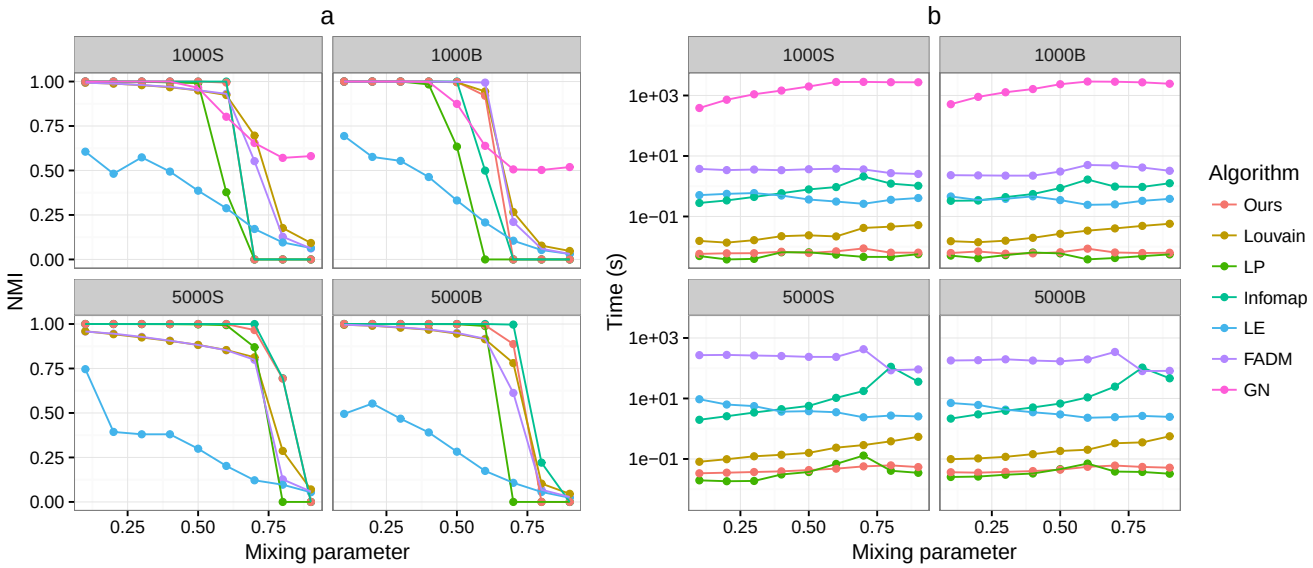

Figure 2. (Colour online) **Tests of the algorithms on the LFR benchmark.** (a) shows the normalized mutual information as a function of the mixing parameter. (b) shows the execution time of different algorithms on the benchmark.

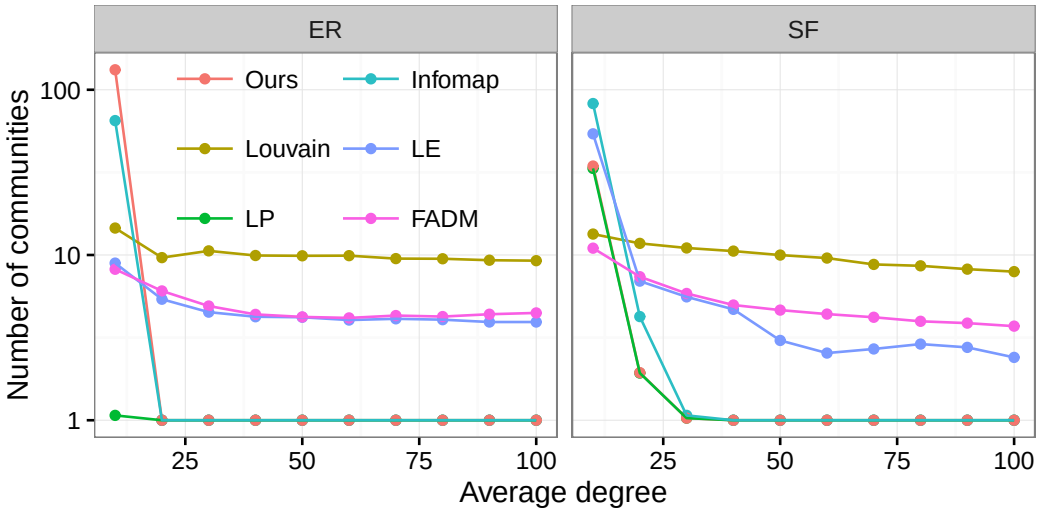

Figure 3. (Colour online) **Tests of the algorithms on ER and SF random networks.** The plots show the number of communities detected by different algorithms as a function of the average degree.

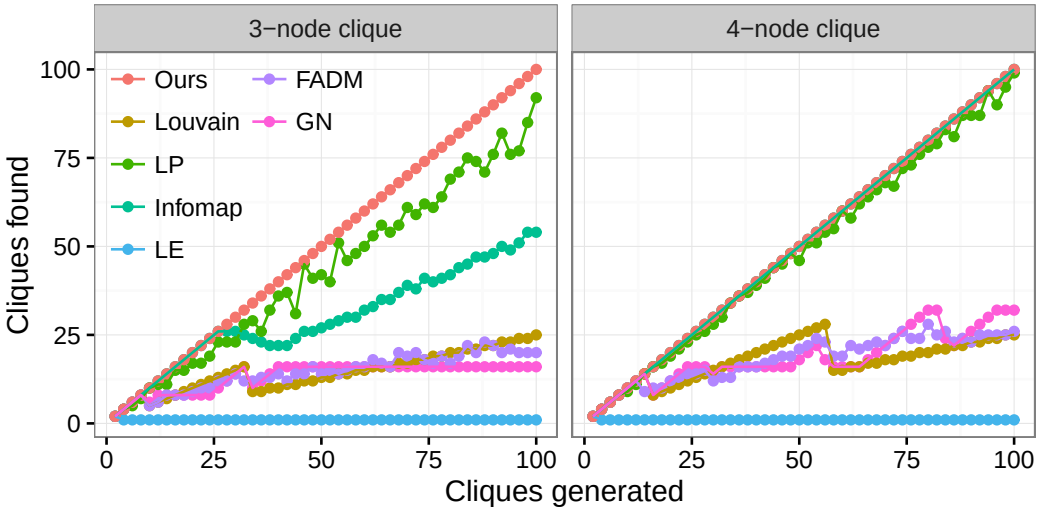

Figure 4. (Colour online) Tests of the algorithms on networks comprised of identical cliques which are connected by single edges.

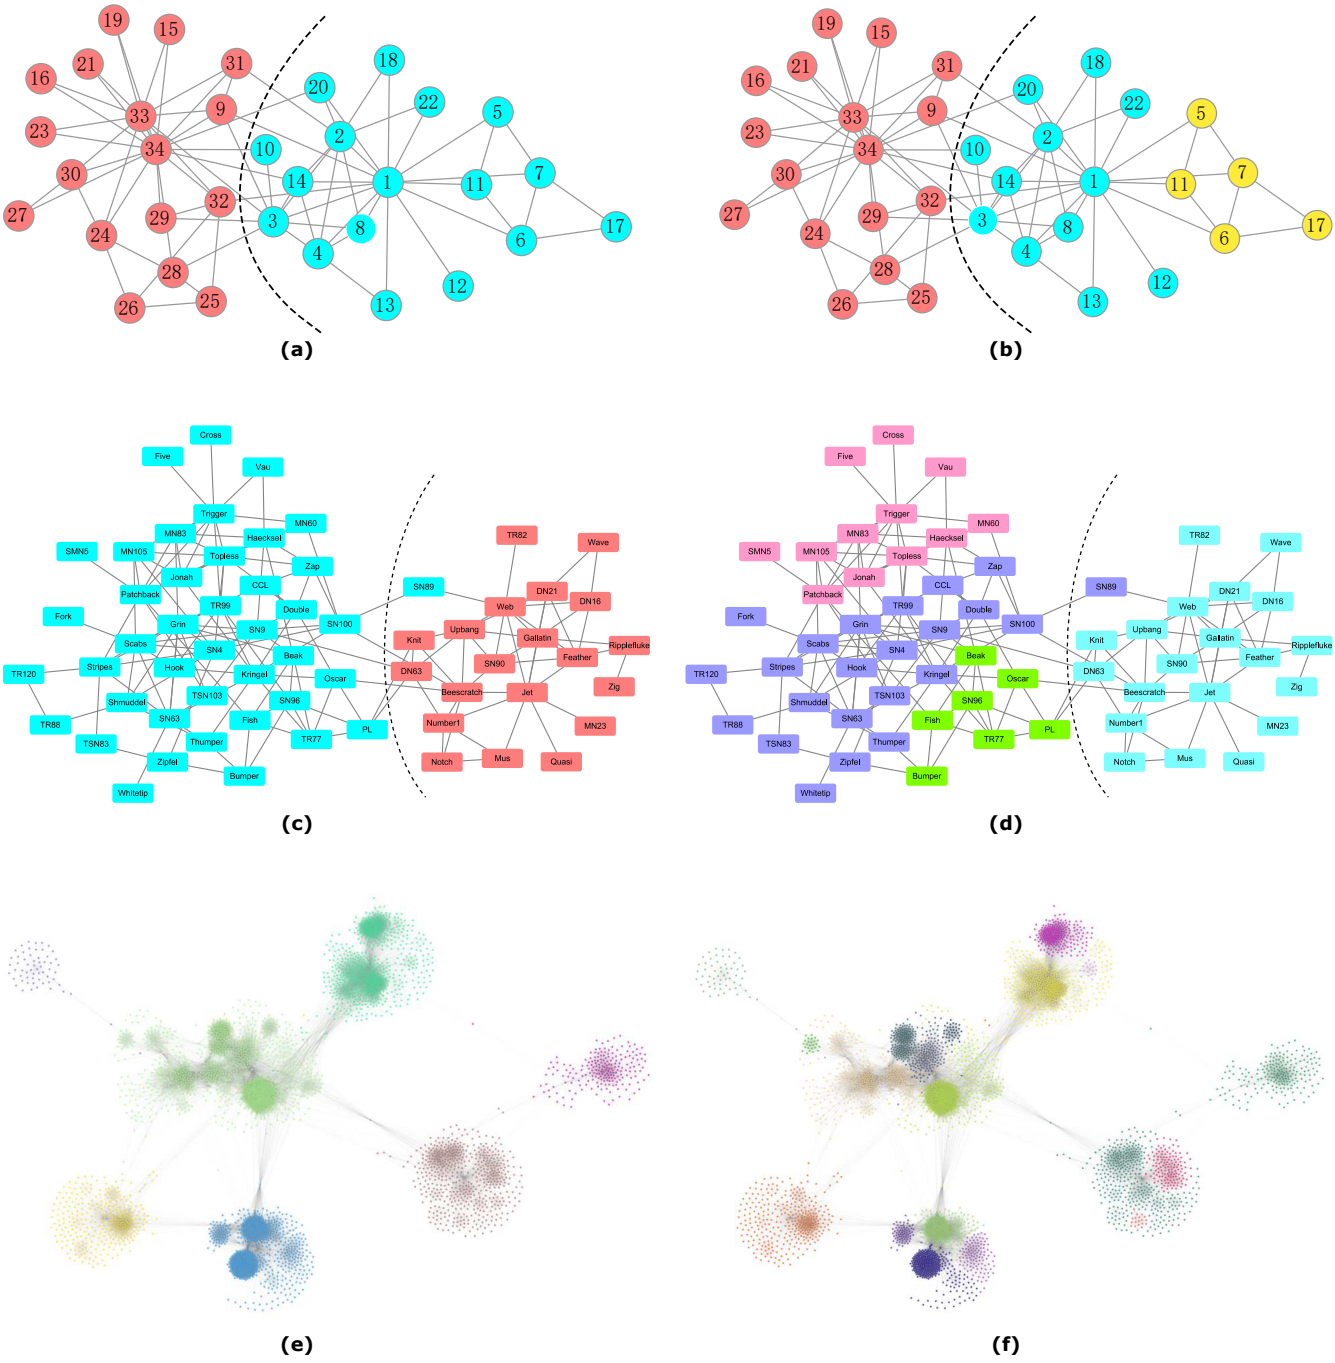

Figure 5. (Colour online) **Communities detected by our method on the real-world networks.** (Real communities are separated by the dashed curve and detected communities are distinguished by different colours). (a) Two communities discovered by our algorithm on the *Karate* network with $\lambda = 0.6$ are identical with the two real ones. (b) Three communities detected by our algorithm on the *Karate* network with $\lambda = 1$. (c) Two communities discovered by our algorithm on the *Dolphins* network with $\lambda = 0.6$, which are identical with the two real ones except for the node 'SN89'. (d) Four communities discovered by our algorithm on the *Dolphins* network with $\lambda = 1$. (e) Seven communities obtained by our algorithm on the *Facebook* network with $\lambda = 0.05$. (f) Nineteen communities obtained by our algorithm on the *Facebook* network with $\lambda = 0.2$.

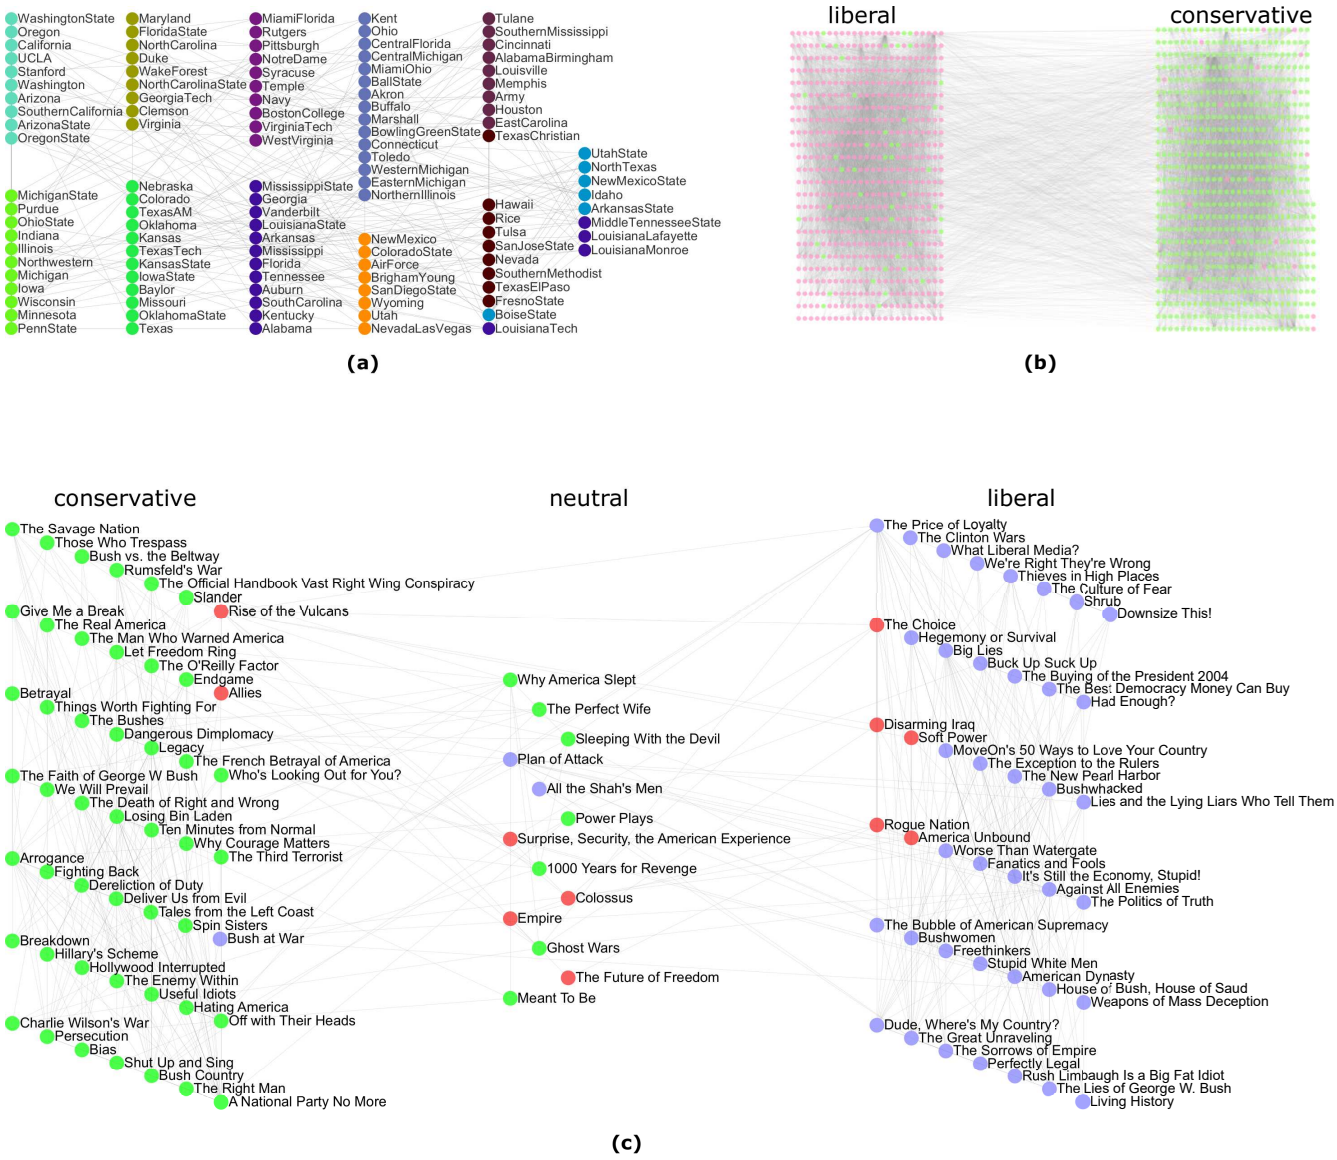

Figure 6. (Colour online) **Communities detected by our method on the real-world networks.** (The real communities are represented by the node positions and the detected communities are distinguished by different colours). (a) The *Football* network and eleven communities obtained by our algorithm on it with $\lambda = 0.6$. (b) The *Polblogs* network: we discovered two large communities using our algorithm with $\lambda = 0.6$. (c) The *Polbooks* network and the three communities detected by our algorithm with $\lambda = 0.6$.

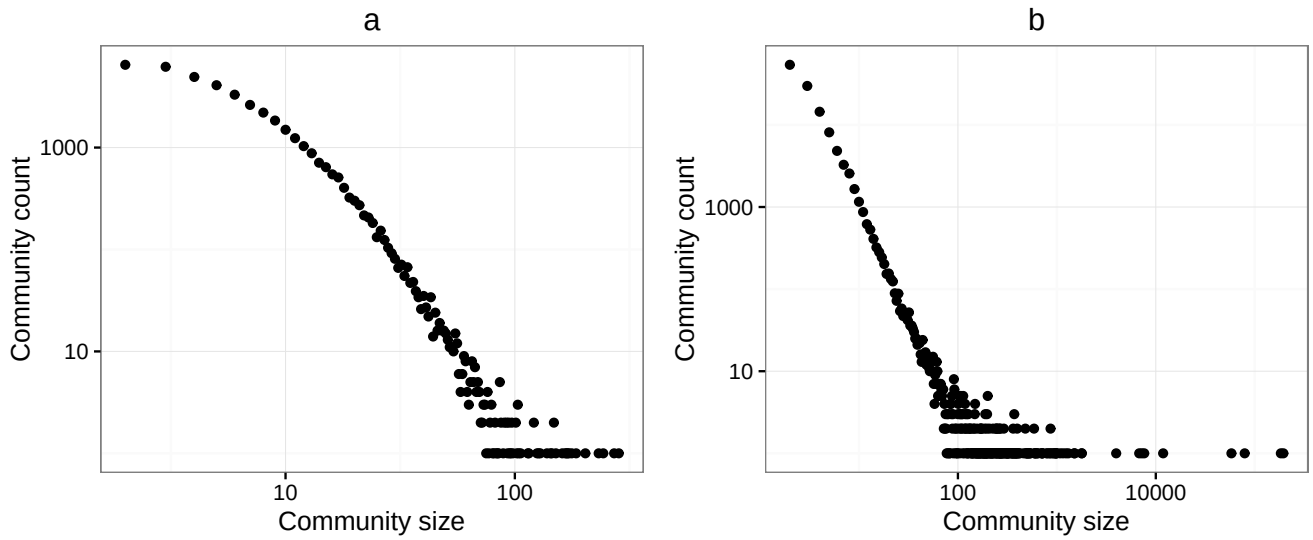

Figure 7. Size distribution of detected communities. (a) and (b) show the size distribution of communities detected by our algorithm with $\lambda = 1$ on Amazon and Youtube networks respectively. In both cases, the behaviour is well reproduced by a power-law, and the exponents are -2.83 and -2.61 respectively.

Author and Ref.	Label	Time complexity
Blondel et al.[16]	Louvain	$O(m)$
Raghavan et al.[17,49]	LP (**L**abel **P**ropagation)	$O(m)$
Rosvall & Bergstrom[29]	Infomap	$O(m)$
Newman[26,49]	LE (**L**eading **E**igenvector)	$O(n^2)$
Treviño et al.[50]	FADM (**F**ast and **A**ccurate **D**etermination of **M**odularity)	$O(n^3)$
Girvan & Newman[12,49]	GN	$O(nm^2)$

Table 1. List of the algorithms used in our experiments. The first column indicates the designers and references, the second one denotes the label used to indicate the algorithm and the last one presents the time complexity of the algorithm. n and m indicate the number of nodes and the number of edges respectively.

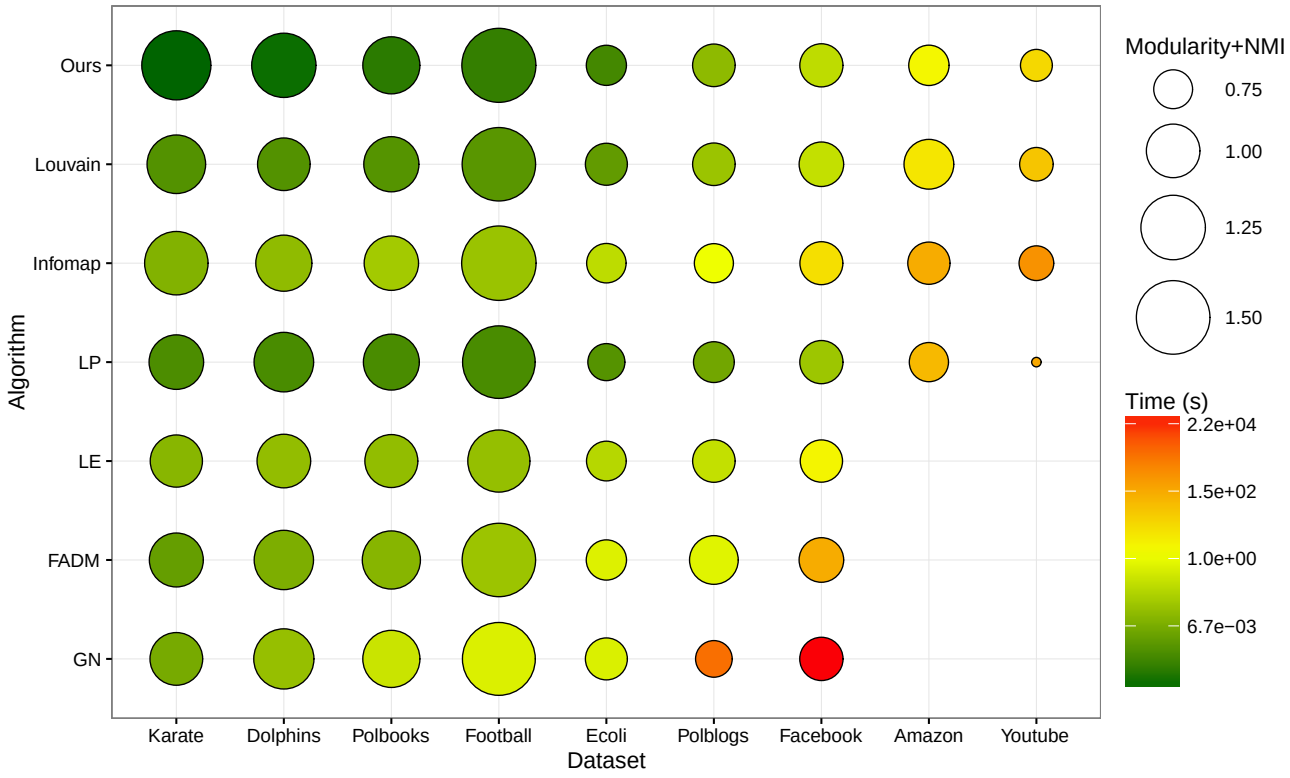

Figure 8. (Colour online) **Tests of the algorithms on real-world networks.** The bubbles show the composite performance and execution time for different algorithms. The composite performance includes two measures: NMI and modularity. For networks without ground truth partitions, we only show the modularity values. For LE, FADM and GN algorithms, the results are presented only for the smaller networks, due to the high complexity of the methods.

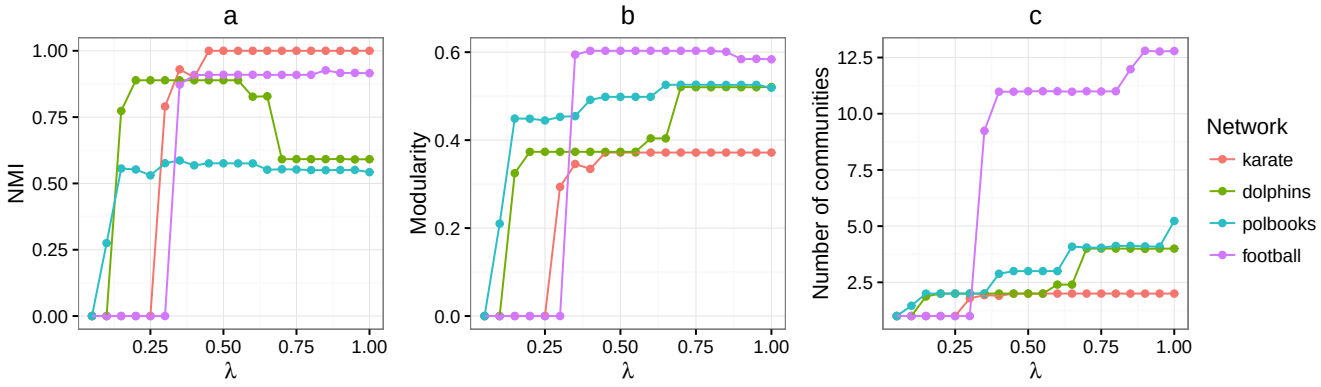

Figure 9. (Colour online) **Resolution parameter analysis on the four real-world networks.** (a), (b) and (c) show the NMI, the modularity, and the number of communities as a function of λ, respectively.

Dataset	Description	Nodes	Edges
Karate[51]	Zachary's karate club	34	78
Dolphins[52]	Dolphin social network	62	159
Polbooks[53]	Books about US politics	105	441
Football[54]	American College football	115	613
E. Coli[55]	Transcriptional regulation data	423	519
Polblogs[56]	Weblogs on US politics (2005)	1490	19090
Facebook[57]	Facebook network	4039	88218
Amazon[58,59]	Amazon co-purchased network	334863	925872
Youtube[58,59]	Youtube friendship network	1134890	2987624

Table 2. A list of real-world networks employed in our experiments.

Index	N_o	N_g	Function clusters	p-value	N_m
1	53	126	cellular respiration	5.20E-89	121
2	6	7	aromatic amino acid family biosynthetic process	3.20E-13	7
3	5	12	aromatic amino acid family biosynthetic process	2.00E-24	12
4	8	23	sulfur metabolic process	1.60E-37	23
5	8	17	iron transport	7.60E-46	17
6	16	18	response to xenobiotic stimulus	7.50E-10	18
7	11	11	stress response	2.10E-21	11
8	15	23	purine biosynthesis	1.20E-31	23
9	10	22	periplasm	1.00E-06	17
10	7	13	signal	5.50E-10	13
11	8	9	DNA binding	8.30E-11	9
12	20	41	Lipopolysaccharide biosynthesis	2.40E-10	27
13	5	6	fatty acid metabolic process	2.70E-09	6
14	15	36	fermentation	1.30E-19	34
15	13	53	flagellum	2.20E-97	41
16	10	22	branched chain family amino acid biosynthetic process	1.50E-14	22
17	7	7	methionine biosynthesis	1.50E-15	7
18	12	16	SOS response	9.90E-28	16
19	5	23	anion transport	5.10E-44	23
20	26	55	oxidoreductase	6.80E-08	55
21	72	144	carbohydrate catabolic process	1.30E-71	143
22	7	10	arginine biosynthesis	8.10E-28	10

Table 3. The function annotations of E. Coli. modules identified by our algorithm. N_o represents the number of operons in each module, N_g represents the number of genes in each module, and N_m represents the number of genes that are matched with the DAVID Escherichia coli database in each module.